

Intelligent Autonomy for Remote Characterization of Hazardous Environments

David J. Bruemmer, Matthew O. Anderson, Julie L. Marble,
Donald D. Dudenhoeffer, Douglas A. Few, Mark D. McKay,

Human Factors, Robotics and Remote Systems Group
Idaho National Engineering and Environmental Laboratory
Idaho Falls, ID 83415

E-mail: {bruedj, matthew, marbjl, dudedd, fewda, mark}@inel.gov

Abstract—Remote characterization of high radiation environments is a pressing application area where robots have the potential to provide benefits in terms of time, cost, safety and quality of data. However, the ability to design robots that can be used effectively has proven to be no easy task. In 2001, the Idaho National Engineering and Environmental Laboratory (INEEL) successfully deployed a teleoperated robotic system coupled with a Gamma Locating and Isotopic Identification Device (RGL&IID) to characterize an area that had been closed to human entry for many years. This paper examines the limitations to the control strategy used and discusses how current efforts at the INEEL are developing intelligent controls that can actively mediate between the human and the robotic elements of the system. The resulting, mixed-initiative control architecture allows the user to shift the level of robot initiative throughout the task as needed. This system offers the opportunity for the human and robot to become a team where each can support the capabilities and limitations of the other.

I. Introduction

As part of the FY 2000 and 2001 Department of Energy Large-Scale Demonstration and Deployment Project (LSDDP), the INEEL collaborated with the Russian Research and Development Institute of Construction Technology (NIKIMT) to develop a novel robotic solution to the problem of characterizing radiation in a remote environment. The resulting Robotic Gamma Locating and Isotopic Identification Device (RGL&IID) integrated DOE Robotics Crosscutting (Rbx) technology with NIKIMT Russian gamma locating and isotopic identification technology [1].

While the new robotic solution offered significant improvements in terms of time, cost, worker exposure and the quality of data acquired, the remote nature of this new technology presented new human-robot interaction challenges. Humans were required to enter the building to instrument the environment with cameras and to assist the robot during the execution of the task. Moreover, during the

actual deployment, the robot was only allowed to move at very slow speeds due to the limitations of visual feedback to the operator.

In response to these challenges, the INEEL has developed a dynamic autonomy control architecture which, unlike the teleoperated approach used in 2001, allows the robot to use its own initiative to support the changing needs and capabilities of the human element. In fact, a truly intelligent robot must be able to do more than decompose high-level commands and make decisions in their absence. We submit that the most interesting and fruitful human robot interaction is when the robot is able to interact with the human as a team member, rather than a tool.

The most important requirement for peer-peer interaction is user trust. For operators to embrace such robots, controls must be crafted to operate and fail predictably. The robots must be able to understand the human's intent and communicate their own understanding of the environment and task. Towards this objective the INEEL has developed novel human-robot interaction (HRI) concepts and interfaces for robust, mixed-initiative interaction between robots and humans. These interfaces utilize simultaneous localization and mapping techniques that capture an abstracted representation of the robot's experience and exploit sensor-suites and fusion algorithms that enhance capabilities for sensing, interpreting, and "understanding" environmental features.

The new mixed-initiative system will remove much of the need for prior instrumentation, remove the need for expert operators, reduce the total number of operators, eliminate the need for human exposure and greatly reduce the time needed for preparation and execution of the task.

II. Teleoperated Approach

The RGL&IID was deployed in July, 2001 at Test Area North (TAN) 616 (see Figure 1). As a result of treating thousands of gallons of liquid nuclear processing waste, there are various levels of contamination present in the facility.

Three rooms within TAN 616 were surveyed using the RGL&IID: the Operating Pump Room, the Control Room, and the Pump Room. All of these rooms are filled with process piping and equipment at various levels, which make a manual survey very difficult and time consuming to perform.



Fig. 1. TAN 616

When compared to baseline assessment methods used by suited radiation control technicians (RCTs), the most significant benefit of the RGL&IID deployment was the quality of the results relative to the safety of the workers. Although the RGL&IID deployment did not eliminate the need for workers to enter the contaminated area, it did reduce the need for human exposure. The RGL&IID was compared to the following baseline activities: the initial RCT entry, an entry to collect video, and a final entry to collect sample information. The RGL&IID was able to collect dose information, video coverage, and isotopes present in a single unmanned entry.

Radiation exposure to workers supporting the RGL&IID deployment was cut by more than a factor of 10 over baseline activities. During baseline characterization, workers received 82mRem of radiation exposure. During the deployment of the RGL&IID, workers received 7mRem of radiation exposure. In addition, the RGL&IID provided radiation survey results instantly and the complete facility survey was accomplished in 3 days. It took workers using baseline characterization methods 3 months to accomplish the same results. The RGL&IID confirmed the presence of Cs-137, Co-60 and Am-241. This data was available within minutes after the RGL&IID performed the scan.

The deployment of the RGL&IID did require more workers than the baseline characterization. However, during the baseline sampling activities, six entries with as many as six individuals per entry were made, totaling 60 work hours spent in the contaminated area. During the RGL&IID demonstration, only two technicians and one RCT were required to enter the contaminated facility for a total of 10 work-hours spent in a contaminated area. All others associated with the project were able to complete the objectives from outside the contaminated areas. As a result of workers spending less time in the radiation areas, individuals involved in the RGL&IID deployment received

10 times less radiation dose than workers involved in baseline activities.

In addition, the two technicians and one RCT who did enter the facility during the demonstration did so only to assist the movement of the RGL&IID up and down a flight of stairs and to check air quality prior to entering the facility. These individuals maintained as much distance between themselves and the highest contaminated areas as possible. In contrast, the baseline samplers were required to come in direct contact with the contaminated material in order to collect representative samples.

The financial cost of collecting the radiation measurements using the RGL&IID was about half the cost of the baseline technology. In addition to the benefit of significant cost reductions, this technology also generates significantly more data. For example, whereas the baseline survey included 10 point samples, the RGL&IID collected about 20 scans. Each scan covers as little as one square foot or as much as several square feet and may have as many as 64 point measurements. Altogether the RGL&IID deployment resulted in over 200 point measurements that covered over 100 square feet of wall and floor area.

Although the 2001 robotic deployment offered a means to reduce human exposure, it did not fully remove the human from the hazardous environment or make it possible for a single human to control the robot. In fact, the baseline survey required only three people, whereas the RGL&IID required six. If robotic systems are to be truly cost-effective and efficient, this ratio of six humans to one robot must be reduced through the use of intelligent control.

Moreover, the data presented above says nothing about the inherent limitations and risks of teleoperation. Teleoperation requires high-fidelity video, reliable, continuous communication, and costly, potentially dangerous efforts to instrument the environment *a priori*. As a mechanical 'subordinate,' the robot was dependent on continuous, low-level input from a human and was poorly equipped to cope with communication failures or changes in operator workload. In fact, while training within a mock-up facility, operators lost control of the vehicle due to a communication failure. Since the last command received by the robot before communications were lost had been a forward velocity command, the robot continued to run through the walls of an adjacent test bed environment.

As is often the case, communication proved to be the limiting factor governing human-robot interaction during the teleoperated deployment. Thick concrete shielding, typical to radiological controls, made it extremely difficult for high-bandwidth communication to support the strictly teleoperated system. As a result, it was necessary for a human to physically place a large antenna directly into the opening of the TAN 616 building. As the robot traveled further from this antenna, the possibility of communication dropouts increased. In fact, operators completely lost contact with the robot at one point during the deployment when the robot traveled out of range. The robot stopped after several seconds once it recognized that

communication had been lost. Since the robot was merely a passive tool, it was unable to reorient itself or attempt to reestablish communication. Fortunately, a human was able to move the antenna slightly further into the doorway of the building and communication was reestablished. Without this good fortune, the robot would have been lost and unable to complete its task.

The 2001 RGL&IID deployment required weeks of preparation including training operators in mock-up environments. Early on, these training exercises indicated that the camera positioned on the robot would not be sufficient to support teleoperation. The camera on the robot could not see the immediate obstacles surrounding the wheels – the very obstacles that posed the greatest threat. As a result, it was necessary to instrument the environment *a priori* with elevated cameras, which were pushed into the environment on wheeled carts. *A priori* placement of tethered cameras is a common practice in nuclear remote inspections throughout the DOE complex. This drawback to teleoperated approaches is further pronounced by the fact that these cameras must be bagged, resulting in additional contaminated waste once the operation is complete.

Although the cameras were deemed sufficient for the task, operators explained that such a strategy is inherently limiting. The first limitation is that adequate lighting is required to support vision-based teleoperation. Secondly, such cameras are usually unable to provide complete visual coverage. In fact, operators reported blind spots when using the same robotic system and cameras within a different, larger building at the site. In one instance, as the robot rounded a corner and left the visual field of one camera, the last thing the operators saw was the robot begin to tip over. Fortunately, the robot did not tip over and the task was completed successfully. Nonetheless, the incident underscores the need for the robot to provide better feedback and, ideally, to be able to take initiative to protect itself in critical situations.

III. Shared Control

Teleoperated systems have often failed to address the limitations of telepresence inherent to current communication technologies. On the other hand, attempts to build and use autonomous systems have failed to acknowledge the inevitable boundaries to what the robot can perceive, understand, and decide apart from human input. Both approaches have failed to build upon the strengths of the robot and the human working as a cohesive unit. Alternatively, mixed-initiative systems can support a spectrum of control levels. Mixed-Initiative robots should:

- Possess intrinsic intelligence and agency.
- Protect humans, environment and self.
- Dynamically shift level of autonomy.
- Accept different modes of human intervention.
- Recognize when help is needed.

Towards these aims, research efforts at the INEEL have developed a novel robotic system that can leverage its own

intelligence to support a spectrum of control levels. We submit that rather than conceive of machines as passive tools or, on the other hand, as totally autonomous entities that act without human intervention, it is more effective to consider the machine as part of a dynamic human-machine team. Within this team, each member is invested with agency – the ability to actively and authoritatively take initiative to accomplish task objectives. For instance, in a remote situation, the robot may be in a better position than the human to react to the local environment, and consequently, the robot may take the leadership role regarding navigation. As leader, the robot can then “veto” dangerous human commands to avoid running into obstacles or tipping itself over. Given the desire to employ robots in hazardous, critical environments, the ability of the human to develop accurate understanding of robot behavior is essential if this capability is to work effectively.

The need for both human and robot to predict and understand one another’s actions presents a daunting challenge. For each level of robot initiative, the user must be able to quickly and accurately predict robotic responses and understand how cumulative robotic actions may converge to fulfill task objectives. Just as the human develops a theory of the robot’s behavior, the robot must be able to understand and predict the behavior of human members of the team. A great deal of future research is necessary to explore the ways in which robots can use direct communications (verbal, gesture, touch, radio communications link) or indirect observation (physically struggling, erratic behavior, unexpected procedural deviation) to infer the need for initiative to be taken. Initiative may also be triggered by the observation of environmental factors (rising radiation levels, the approach of additional humans, etc.). The robot’s expectations must allow it to recognize human limitations and anticipate human needs without second-guessing the human’s every move. When robots do intervene, the human’s understanding of robot behavior must be able to explain why the robot has stepped in and what this shift in control means for the task at hand.

The benefits of allowing the team members to change roles within the team significantly increases team flexibility and reliability in task performance. However, if the interface and human-robot system are not designed in accordance with critical principles of human factors in mind, dynamic role changing may result in mode confusion, loss of operator situation awareness, loss of operator confidence in assuming supervisory control, and degraded and potentially catastrophic performance [2]. Systematic human-centered design is necessary to insure that the robot autonomy conforms to the ways in which humans assign and manage tasks.

Appropriate feedback is required when roles and levels of initiative change. Failure to inform the operator when the robot has overridden commands will lead to distrust of the system, unless the behavior is beneath the level of operator concern. Feedback from the robot should not only include

the mode change, but also an indication of the reason for the change. For optimal performance of the team, the human must be able to develop expectations regarding when and why the robot will be motivated to change its level of autonomy. In order for the human to comprehend and exploit robot initiative, the intelligent control system should be structured hierarchically such that each mode of human intervention is accompanied by a well-understood shift in the initiative afforded to the robot.

IV. System Design

The resulting robotics system can interpret and fuse a variety of range sensor information including inertial sensors, compass, wheel encoders, laser range finders, computer vision, thermal camera, infrared break beams, tilt sensors, bump sensors, sonar, and others. The robot does not assume that these sensors are working correctly, but rather continuously evaluates its own perceptual capabilities and behavior. The robot is able to abstract information about the environment at many levels including terse textual descriptions of the robot's local surroundings and the choices (depending on the level of autonomy) which face the human user.

Our research to date has developed a control architecture that supports the following modes of remote intervention: *Teleoperation*, *Safe Mode*, *Shared Control*, and *Full Autonomy*. Our *teleoperation* mode is based on the interaction substrate used in previous INEEL teleoperated robotic systems. With feedback from those who have used this system, we have added additional motion control methods including several different kinds of joysticks as well as keyboard and touch screen options. Within teleoperation mode, the user has full, continuous control of the robot at a low level. The robot takes no initiative except to stop once it recognizes that communications have failed.

In *safe mode*, the user directs the movements of the robot, but the robot takes initiative to protect itself. The robot will stop its motion just before a collision. The robot also continuously assesses its own state and the validity of its diverse sensor readings and communication capabilities. The robot will refuse to undertake a task if it does not have the ability (i.e., sufficient power or perceptual resources) to safely accomplish it.

In *shared control* mode, the robot takes the initiative to choose its own path, responds autonomously to the environment, and works to accomplish local objectives. Although the robot handles the low level navigation and obstacle avoidance, the user supplies intermittent input, often at the robot's request, to guide the robot in general directions. The problem of deciding how and when the robot should ask for help has been a major line of HRI enquiry in our human subject experiments.

In the *fully autonomous* mode, the robot selects its own routes, accepting no user input except high-level tasking such as "follow that target" or "search this area."

For each of these levels of autonomy, perceptual data is fused into a specialized interface (shown in Figure 3) that

provides the user with abstracted graphical and textual representations of the environment and task appropriate for the current mode. Immediate obstacles that inhibit motion are shown as red ovals and resistance to motion is shown with arcs emanating from the wheels. The robot relays a great deal of synthesized, high-level information (including suggestions and requests for help) to the user in a textual form using the feedback textbox within the image window.

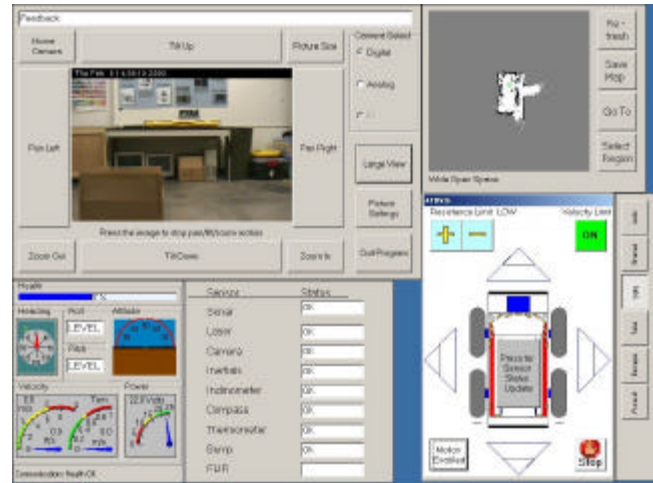


Fig. 3. Current Interface

Also note that the robot provides textual reports on environmental features at the bottom of the map window and reports on communications status at the bottom of the robot status window. The robot status window provides a variety of information including pitch and roll, power, heading, speed and a fusion of this information into a single measurement of "health." The user can shift into shared or teleoperation mode and then move the robot by touching the arrows or using a joystick. Also, it is possible to pan and tilt the camera by touching regions of the visual image.

The fundamental aspect of a human team that distinguishes it from a simple group is the presence of a shared goal. We must have some means to represent this common goal with a common form of representation that is meaningful to all members. Effective teams typically cooperate and anticipate the needs of teammates via a shared mental model of the task and current situation [3]. If we want humans, air vehicles and ground vehicles to work as a team, we need to develop an appropriate level of discourse, including a shared vocabulary and a shared cognitive workspace, collaboratively constructed and updated through interaction with the real world.

We have chosen to address this need by building a map that consists of terrain overlaid with semantic abstractions generated through autonomous or user-assisted recognition of environmental features. The current mapping algorithm is based on simultaneous localization and mapping (SLAM) work done at the Naval Research Laboratory [4]. At the present time, we have successfully implemented the

algorithm on individual robots and integrated it with the control architecture and communication protocol.

This real-time semantic map, constructed collaboratively by human and machine, serves as the basis for a spectrum of mutual human-robot interactions including tasking, situation awareness, human-assisted perception and collaborative environmental “understanding.” Collaborative construction of the map enhances each individual team member’s understanding of the environment and provides a shared semantic lexicon for communication. For the user, the map provides point-and-click user validation and iconographic insertion of map entities. The user can verify or remove entities, which have been autonomously added and can add new entities, which the robot was unable to find. The robot can use the workspace to communicate about the task and environment both graphically (e.g. “The highlighted area has been searched!”) and verbally (“Landmine found near Victim 2!”) using the semantic names which have been assigned within the shared cognitive workspace. Conversely, the human can task the robot in much the same way.

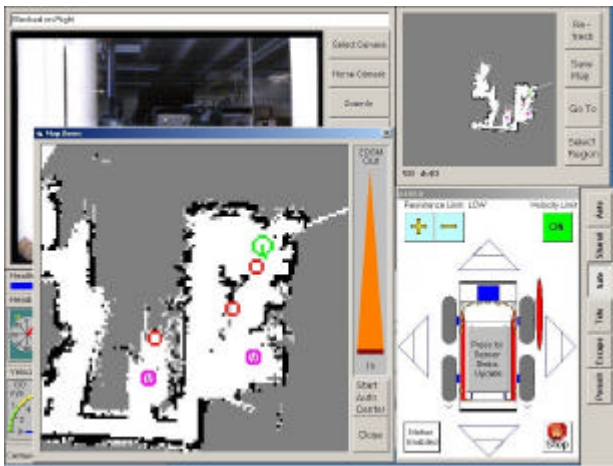


Figure 4: Collaborative Cognitive Workspace

V. Conclusions

Intelligent controls that permit a spectrum of operator intervention can greatly improve on the opportunities provided to the operators of a strictly teleoperated system such as the one used in the RGL&IID deployment. The human user can switch between these modes to cope with different components of the task. For instance, when a user wishes to move into a new room s/he simply points the robot at a door and then allows the robot to guide itself through the doorway – a task that can take teleoperators many minutes of trial and error.

The robot is often able to make better judgments about its environment (i.e., local navigation) than distal human controllers. Consequently, we have created modes of control where the robot monitors human command input and infers the need to supplement or override human action. The robot has the power to refuse to undertake commands

from the user that are deemed by the robot to pose a threat to itself or its environment. This engenders a host of new questions regarding how to interleave human control and robotic initiative.

For a robotic system to gracefully accept a full spectrum of intervention possibilities, interaction issues cannot be handled merely as augmentations to a control system. Instead, opportunities for operator intervention must be incorporated as an integral part of the robot’s intrinsic intelligence. The robot must be imbued with the ability to accept different levels and frequencies of intervention. Intelligent controls require not only that the robot possess intrinsic intelligence, but also that the control paradigm be intelligently structured to support both the human and the machine. A control paradigm that is not designed to be flexible will impede even the most intelligent robots and the most skilled operators.

References

1. Anderson, M., Conner, C., Daniel, V., McKay, M., Yancey, N., “Demonstration of the Robotic Gamma Locating and Isotopic Identification Device,” In *Proceedings from the American Nuclear Society Spectrum 2002*, Reno, NV, August 2002.
2. Yen, J., Yin, J., Ioerger, T., Miller, M., Xu, D., & Volz, R. “CAST: Collaborative Agents for Simulating Teamwork,” In *Proc. of the 17th Int. Joint Conf. on Artificial Intelligence*, Seattle, WA, 2001.
3. Casey, S. (1998). *Set Phasers on Stun: And Other True Tales of Design, Technology, and Human Error* (2nd edition). Santa Barbara, CA: Aegean Publishing Co.
4. Yamauchi, B., Schultz, S., Adams, William. (1999). “Integrating Exploration and Localization for Mobile Robots.” *Adaptive Behavior* 7(2), 217-229.